

ABSTRACT OF THE DISCLOSURE

A storage system implements a storage operating system configured to concurrently perform speculative readahead for a plurality of different read streams. Unlike previous implementations, the operating system manages a separate set of readahead metadata for each of the plurality of read streams. Consequently, the operating system can “match” a received client read request with a corresponding read stream, then perform readahead operations for the request in accordance with the read stream’s associated set of metadata. Because received client read requests are matched to their corresponding read streams on a request-by-request basis, the operating system can concurrently perform readahead operations for multiple read streams, regardless of whether the read streams’ file read requests are received by the storage system in sequential, nearly-sequential or random orders. Further, the operating system can concurrently perform speculative readahead for the plurality of different read streams, even when the read streams employ different readahead algorithms. The invention may be implemented by file-based or block-based storage systems, or combinations thereof.